

(12) UK Patent Application (19) GB (11) 2 298 111 (13) A

(43) Date of A Publication 21.08.1996

(21) Application No 9501834.7

(22) Date of Filing 31.01.1995

(71) Applicant(s)

Videologic Limited

(Incorporated in the United Kingdom)

Unit 8, Home Park Industrial Estate, Kings Langley,
Hertfordshire, WD4 8LZ, United Kingdom

(72) Inventor(s)

Martin Ashton

Simon James Fenny

Hossein Yassale

(74) Agent and/or Address for Service

Reddie & Grose

16 Theobalds Road, LONDON, WC1X 8PL,
United Kingdom

(51) INT CL⁶

G06T 15/50

(52) UK CL (Edition O)

H4T TBBA

(56) Documents Cited

GB 2246057 A EP 0613099 A1 EP 0610004 A2

US 5305430 A

Des. and Sim. of Opera Lighting and Projection Effects
Dorsey Sillon Greenburg pgs 41-50 Com. Graphics

(58) Field of Search

UK CL (Edition N) H4T TBAD TBBA TBBX TBEA TBEC

INT CL⁶ G06T 1/20 1/60 15/00 15/10 15/30 15/40

15/50 15/60

Online: WPI; H4T Library




(54) IMPROVEMENTS RELATING TO COMPUTER 3D RENDERING SYSTEMS

(57) An apparatus for shading a three dimensional image for display on a screen operates by first supplying data defining a group of infinite surfaces representing each object in the image. A depth value is generated for each surface for each elementary area of the display in dependence on the distance of that surface from an image plane. The apparatus then determines whether or not any surface is visible at that elementary area. One or more objects may be designated as a light volume. A depth value is determined for each surface of the light volume for each elementary area of the display in dependence on distance of that surface from an image plane. The elementary area of the display is then shaded in dependence on the object surface visible at that elementary area and its location in relation to the light volume.

To facilitate easier processing of the shading the display may be sub-divided into a plurality of sub-regions. The size of the sub-regions is dependent on the complexity of the image of each area of the screen. Preferably the determination of the depth values for surfaces for groups of pixels in one sub-region is interleaved with the determination of depth values for surfaces for groups of pixels in other sub-regions.

The apparatus performing the processing comprises a plurality of groups of processing elements for determining depth values for surfaces in the image for each elementary area of the display. Each group of processing elements has an associated cache memory means for storing defining the various surfaces in the image. The groups of processing elements are coupled together by a bus. Data read from the associated memory of the first group of processing elements is transmitted via the bus to other groups of processing elements for the determination of depth values until all data from the first associated memory means has been read. Reading of data then takes place from the associated memory means from a second group of processing elements and so on.

GB 2 298 111 A



2 / 6

FIG. 2A.




FIG. 2B.





FIG. 2C.






FIG. 5




FIG. 6




FIG. 7




FIG. 8

FIG. 9





FIG. 11







FIG. 10



IMPROVEMENTS RELATING TO COMPUTER 3D RENDERING SYSTEMS

This invention relates to computer 3D rendering systems of the type described in our British patent application No. 9414834.3.

5 British patent application No. 9414834.3 describes a 3D rendering system in which each object in a scene to be viewed is defined as a set of infinite surfaces. Each elementary area of a screen on which an image is to be displayed has a ray projected through it
10 from a viewpoint into the three-dimensional scene. The location of the intersection of the projected ray with each surface is then determined. From these intersections it is then possible to determine whether any new intersected surface is visible at that elementary
15 area. The elementary area is then shaded for display in dependence on the result of the determination.

 The system is implemented in a pipeline type processor comprising a number of cells each of which can perform an intersection calculation with a surface. Thus
20 a large number of surface intersections can be computed simultaneously. Each cell is loaded with a set of coefficients defining the surface for which it is to perform the intersection test.

 The system is capable of generating shadows by
25 defining volumes in shadow in the same way as solid objects, that is to say they are defined as a set of infinite surfaces. These shadow objects are then processed in the same way as solid objects. Once all the shadow and non-shadow objects have been processed, it is a
30 straightforward matter to determine whether the surface visible at a particular pixel is in shadow by comparing the position of the visible surface with the furthest forward facing surface and the nearest reverse facing surface of the shadow object and checking to see whether
35 the visible surface is within that range. If it is then a flag associated with that surface is set to indicate that

a modification to its shade or colour needs to be made to simulate the effect of shadowing.

Preferred embodiments of the present invention seek to improve the performance of the pipeline processor
5 described in British patent application No. 9414834.3 by sub-dividing a scene into a variety of sub-regions for processing. These sub-regions are preferably rectangular tiles.

Another embodiment of this invention seeks to
10 provide a system in which spotlight-type effects can be generated within a scene.

The invention is defined with more precision in the appended claims to which reference should now be made.

Preferred embodiments of the invention will now
15 be described in detail, by way of example, with reference to the drawings, in which:-

Figure 1 schematically shows a spotlight type effect to be modelled;

Figures 2A, 2B, and 2C show the way in which
20 concave objects are dealt with by the system;

Figure 3 shows schematically a rectangular tile containing four objects to be processed by the system;

Figure 4 shows schematically the same tile
25 divided into four sub-tiles for processing;

Figure 5 shows further sub-divisions of tiles;

Figure 6 shows a variable tile size arrangement;

Figure 7 shows schematically the arrangement
30 for processing elements of the prior art system;

Figure 8 shows an alternative arrangement for the prior art system;

Figure 9 shows an arrangement of tiles having a different number of surfaces per tile for mixed
35 processing in accordance with an embodiment of the invention;

Figure 10 shows a block diagram of a pipeline processor embodying the invention;

Figure 11 shows the manner in which surface data is retrieved from memory for use in an embodiment of the invention; and

5 Figure 12 shows a plurality of pipeline processing embodying another aspect of the invention.

10 In order to make computer-generated three-dimensional images more lifelike, it is important to be able to simulate the effect of a spotlight in a night time scene. For example, car headlights, search lights, and landing lights are all things which it may be desired to simulate. The type of effect is illustrated in Figure 1 in which a light source 2 illuminates a cone-shaped volume 4 of light which is projected onto a surface 6. An object 8 positioned on the surface is illuminated by the
15 cone of light 4.

20 The cone-shaped volume 4 is referred to as a light volume. This is to all intents and purposes, the opposite of the shadow volumes which are described in British patent No. 9414834.3. A light volume is a region of space whose shape is determined by the characteristics of the light source. All objects or parts of objects which are outside this area are in shadow with respect to the light source.

25 In order to render the scene the object 8 is defined as a set of surfaces as is the light volume 4. In the case where three-dimensional shapes are concave it is preferable, as in British patent application No. 9414834.3 to model the object as a collection of convex objects. This is illustrated in Figures 2A, 2B, and 2C. Figures 2A and 2B show two concave objects arranged to intersect to produce an object with concavities. Figure 2C shows how an object with concavities can be split into two convex objects.

30 As described in British patent application No. 9414834.3, after all the non-shadow convex objects have been processed, we have, at any particular pixel, the distance to the visible surface at that point. Thus, to find whether the surface at that pixel is in a light object we then need to process the surfaces defining the
40 light volume 4 to find the furthest forward facing surface

and the nearest reverse facing surface of the light object. If the furthest forward facing surface is closer to the viewpoint than the nearest reverse facing surface and the visible surface between them, then that visible surface is within the light volume. If the furthest forward surface is nearer to the viewpoint and the nearest reverse surface and the visible surface is between them then the surface of the object is within the light volume and a flag is associated with it to indicate that its colour will need to be modified to simulate the effect of light. If the furthest forward surface is further from the viewpoint than the nearest reverse surface then the object is outside the light volume and a flag has to be set to indicate that this colour will need to be modified to simulate the effect of shadowing.

Light volumes are processed before the shadow volumes of GB 9414834.3. If a surface of a pixel is outside the light volume then the shadow flag associated with the surface is asserted. In the case of surfaces falling within the light volume shadow volumes are generated within the light volume to simulate shadows cast by the light source and hence objects within the light volume cast shadows.

Light volumes do not require any additions to the architecture as given in Patent No. 9414834.3. The only change that is required is that the operation code associated with each plane has to be extended so that planes can be classed as light volume objects, just as there is currently an operation code to indicate shadow volume objects. The light volume object uses the same logic as the shadow volume object except that the output to the shadow flag is inverted if the surface is a light volume surface.

Turning now to ways in which the performance of the rendering system may be improved, it has been proposed that the screen be divided into a number of sub-regions or tiles. Then, for each tile, only those objects within the tile will be processed. This, decreases the average number of surfaces to be processed when determining which surface is visible at any particular pixel.

A gross elimination process such as surrounding complex objects with bounding volumes can be first used to test which objects apply to the current region being rendered. The bounding volume is defined by finding the areas to which particular objects apply. The sub-regions can, in general, be any shape. However, the simplest technique is to divide the screen into rectangular tiles.

All the bounding volumes of objects are projected onto screen space (the array of pixels defining the image plane) and are tested against corners of the tiles. Those bounding volumes of objects which are completely outside the tile are discarded for that tile. Thus, the number of surfaces which need to be processed per pixel within a tile becomes less and hence the total time to render an image is reduced since the total processing time for all the tiles will be reduced.

Figure 3 shows the bounding volumes of four objects projected onto an array of twenty pixels by twenty pixels. If each object is defined by ten surfaces, then the total time to render Figure 3 will be 400 pixels x 40 surfaces which equals 16,000 cycles. Figure 4, however, shows an area of 20 x 20 pixels sub-divided into four sub-arrays of 10 x 10 pixels. Thus the time to render the objects in Figure 4 will be defined by the time to render the object in end 10 x 10 pixel tile. This will be 100 pixels by 10 surfaces which equals 1,000 cycles. Thus to render the whole 20 x 20 array, the time take will be 4,000 cycles. Thus there is a 75% saving in processing time using this arrangement.

However, the objects of the image are not normally evenly distributed over the entire screen. Therefore, using tiles with variable sizes allows the same performance enhancement to be gained with fewer tiles or better performance with the same number of tiles. This is shown in Figures 5 and 6.

In Figure 5 there are a total of seven objects and 16 tiles each 5 pixels by 5 pixels. Therefore, the time taken to render a scene will be 16 x 25 x 10 which gives 4,000 cycles.

In Figure 4 the number of tiles is reduced by having three 10 pixel x 10 pixel tiles encompassing the larger objects and for 5 pixel x 5 pixel tiles encompassing the smaller objects. The time to process the 10 x 10 pixel tiles will be 100 pixels x 10 surfaces x 3 tiles which equals 3,000 cycles. The time to process the 5 pixel x 5 pixel tiles is 25 pixels x 10 surfaces x 4 tiles which equals 1,000 cycles. Thus the total time to render Figure 6 will be 4,000 cycles, the same as it would take to render Figure 5. However, Figure 6 only has 7 tiles as opposed to the 16 of Figure 5. Thus, a saving is made on the sub-division of the scene into tiles.

In a preferred embodiment of the invention projection techniques are used which project boundary boxes around complex objects. Firstly the distribution of objects around the visible screen is determined and this then allows suitable tile sizes to be defined in dependence on that distribution.

The surfaces defining the various objects are stored in one contiguous list as described in Patent No. 9414834.3. However, for each tile, (a sub-divided area of the screen), there is another list of object pointers which point to the individual objects which reside in that tile. This is to avoid storing identical surfaces for each tile, as one object made of many surfaces could be in a number of tiles. This is illustrated in Figure 11. At the start of the object pointer list there is a x,y start position and a x,y size. This gives the top left hand position and the width and height in pixels of the tile. Thus, allowing for variable tile sizes.

At the start of rendering the pointer list is traversed and the x,y position and size is read. The x,y position is stored in register 44 of Figure 10. This gives the starting position for the rendering. As each pointer is read the appropriate surfaces pointed to are also read and stored in the internal memory. After all the pointers have been read the internal memory contains all the objects to be processed. The objects are processed as normal. After the tile has been rendered

the next object pointer list is traversed as described earlier and the internal memory contains all the objects for the next tile. This continues until all the tiles have been completed.

5 A further improvement to sub-dividing the screen into sub-regions (tiles) is achieved by having the ability to process the tiles non-sequentially.

 The system described in British patent application No. 9414834.3 uses multiple processing
10 elements that operate on a group of adjacent pixels at any one time. For example, if a tile consists of $N \times M$ pixels there may be an N element parallel processing engine which will process one row of the tile at a time for each input plane. The processing time, per row, in
15 terms of clock cycles will be equal to the number of planes in or associated with that tile. Such an arrangement is shown in Figure 7. If there are tiles which have a very small number of planes, then the parallel processing elements may have to stay idle while
20 the last processed row of pixels is post-processed (e.g. has texture and shading added) by a post processing system 20. This receives pixel data from the N processing elements 22, via an N stage buffer 24. The speed will be limited by the speed of the post-processing unit 20.

25 One way to overcome this speed inefficiency is to ensure that the post-processing unit 20 can match the peak output rate of the N processing elements 22 for a small numbers of surfaces. This would make the post-processing unit more complex and expensive. Another
30 solution would be to interpose a frame buffer 26 between the N stage buffer 24 and the post-processing unit 20 as shown in Figure 8. Preferably this frame buffer would accommodate the same number of pixels as the display. This, however, would again lead to additional expense.

35 A preferred embodiment of the present invention solves this problem in an efficient and inexpensive manner by intermixing the processing of rows from tiles so that rows from tiles with a smaller number of planes contained within them are interleaved with rows from tiles with a
40 large number of planes.

For example, Figure 9 shows tiles A, B, C and D all having different numbers of surfaces defining objects within them. If tiles are processed one at a time in the A, B, C, D order then during the processing of the first
5 two tiles the N element parallel processing array 22 will spend idle time since only a small number of planes need to be processed. However, by processing a row from tile A followed by a row from tile C which has a large number of planes associated with it, then the average output rate
10 of data from the N processing elements 22 will be reduced. Thus, the only buffering needed will be the N stage buffer shown in Figure 7.

The surfaces are stored exactly as detailed above. However, rather than reading in a list of objects
15 for just one tile into the internal memory, two adjacent object lists are read and stored in the internal memory. The first set of surfaces are processed for one line of the first tile which is then followed by processing the second set of surfaces for the second tile. The x,y
20 position of the tiles are stored in two registers such as numbered 42 in Figure 10. These are swapped between as lines from the two different tiles are processed. The tiles which are to be processed together have to have their object lists adjacent in the external memory. In
25 the example above with reference to Figure 11, tile 1 would be tile A and tile 2 would be tile C.

Figure 10 shows an implementation of a system of the type described in British patent application
No. 9414834.3 based on an array of pipeline processing
30 cells 40, a precalculation unit 40,46,50,48 and a memory 42. The memory stores a list of instruction words comprising surface parameters (ABC), control bits and surface attributes. The practical and cost effective VLSI implementation of such a system requires that
35 instruction words would be stored in an onchip cache memory. The reasons for this are two-fold. First, the instruction words are in a very long format (typically 96 bits) and therefore sharing external memory with, for example a 32 bit processor would involve multiple bus
40 cycles per instruction word fetch. Secondly, the rate

of dispatch of instruction words is directly proportional to the surface processing performance of the system. Taken together these points indicate that any external memory solution must use high speed, wide and dedicated memory.

Using a cache implementation of this system exploits the fact that during the processing of a given scene many references are made to the same instruction word for processing a particular surface. In the variable tile size enhancement presented above the tile size can be optimised such that all tiles contain a roughly constant number of surfaces and thus a roughly constant number of instruction words. If the cache size is larger than this value then all instruction words for a given tile will fit in the cache. Under these conditions the advantage in time is given by the following:

Time to process a tile from external memory =
 $PC \times I \times MEXT$

Time to process a tile from cache memory = $I \times (MEXT + (PC - 1) \times MINT)$

Percentage improvement = $1 - (MEXT + (PC - 1) \times MINT) / (PC \times MEXT)$ where T = tile size, PE = the number of pipeline elements, PC = the number of pipeline cycles = T/PE , I = the number of instruction words per tile, $MEXT$ = external memory access time, $MINT$ = cache memory access time.

As scene complexity increases practical limits to cache size dictate that not all instruction words in the given tile can be cached.

The problem of scalability of memory capacity in a multichip implementation of the system described in patent application 9414834.3 can be solved by a further embodiment of the invention.

An implementation of a 3D rendering system consists of N 3D processing units (henceforth referred to as ISP), a system processor and system memory, is shown in Figure 12. Each ISP contains an M instruction word memory which functions as described above. In such a configuration all ISP's are working on a different output

line of the same tile, and are thus all operating on the same instruction word at any given time. For the purpose of this description the ISP working on tile line n is referred to as master, the ISP working on tile line $(n+1)$ is referred to as slave_1, that working on tile line $(n+m)$ as slave_m etc. The memory controllers of all ISP's are connected by a memory interconnect bus 50 consisting of the following elements:

- 1) Data Bus 50: Width of instruction word: B_i
directional
- 2) Load In / Load Out 54: Synchronous daisy chain signals which indicate transfer of cache masterchip :
Input / Output.
- 3) Terminal Word 56: Indicates that current instruction word is the last in the list - driven by chip which is currently transmitting data.
- 4) Data Valid 58: Driven by the current transmitter to qualify data on data bus.

The master ISP 52 is designated by having its load_in input tied '1' at reset. The master is responsible for following a hierarchical list of pointers which ultimately reference instruction words. At the start of a given tile a memory load phase begins in which the master ISP will begin fetching instruction words from system memory 60, assembling the words internally and transmitting them in a synchronised fashion to three destinations : the master internal memory 68, the master processing pipeline and the memory data bus 50. At this point only the master is storing instructions in its memory. The data valid signal is necessary in systems where the master may be denied access to system memory for periods of time. It is also used to indicate a data bus turn around cycle which occurs when memory data bus ownership changes.

On the transfer cycle when the master memory is full, the master drives load_in of slave_1 to '1', indicating that slave_1 must begin storing instruction words in its memory. This process continues down the chain of slave ISP's until the final instruction word is passed out by the master. At this point terminal_word is driven to '1', and the slave which is currently loading its memory latches the value of the terminal memory address. This completes the memory load phase.

There follows a series of memory read phases which continue until the current tile processing is complete. Memory read commences with the master ISP reading out its instruction word list and passing the values to its processing pipeline and to the memory data bus. On the final access, the master drives load_in of slave_1 to '1', indicating that slave_1 must begin reading instruction words from its memory. This process continues down the chain of slave ISP's until the final instruction word is passed out by the final slave. At this point terminal_word is driven to '1', and the process repeats from the master. Note that there will be a lost cycle between each change of data bus mastership to allow for the turnaround time of the buffers.

CLAIMS

1. A method for shading a three dimensional image for display on a screen comprising the steps of representing each objection in the image as a group of surfaces, for each elementary area of the display
5 projecting a ray into the 3-dimensional image, determining the location of the intersection of the ray with each surface, determining whether any intersected surface is visible at that elementary area, and including the step of
10 designating an object as a light volume, determining the intersection of each ray with the light volume, and shading an elementary area in dependence on the surface visible at that elementary area and its location in relation to the light volume.
- 15 2. A method for shading a three dimensional image for display on a screen comprising the steps of representing each objection in the image as a group of surfaces, determining a bounding volume for each object projected onto the display screen, subdividing the display
20 into a plurality of sub-regions, the size of the sub-regions being dependent on the complexity of the image in each area of the screen, and for each pixel of each sub-region projecting a ray into the image, determining the location of the intersection of the ray with each surface
25 of each object contained within that sub-region, determining whether any intersected surface is visible at that pixel and shading the pixel in dependence on the result of the determination.
- 30 3. A method for shading according to claim 2 including the step of interleaving the determination of ray surface intersections for rows of pixels in a sub-region with the determination of ray/surface intersections for rows of pixels in other sub-regions.

4. A method for shading a three dimensional image for display on a screen comprising the steps of representing each object in the image as a group of surfaces, subdividing the screen into a plurality of sub-
5 regions, for each pixel of each sub-region projecting a ray into the image, determining the location of the intersection of the ray with each surface of each object contained within that sub-region, determining whether any intersected surface is visible at that pixel, and shading
10 the pixel for display in dependence on the result of the determination, wherein the determination of ray/surface intersections for groups of pixels in one sub-region is interleaved with the determination of ray/surface intersections for groups of pixels in other sub-regions.

15 5. A method for shading according to claim 4 wherein the groups of pixels comprise lines of pixels in the sub-regions.

6. A method for shading according to claim 4 or 5 wherein groups of pixels from sub-regions in less complex
20 parts of the image are interleaved with groups of pixels from sub-regions in more complex parts of the image.

7. Apparatus for shading a three dimensional image for display on a screen comprising a plurality of groups of processing elements for determining depth values for
25 surfaces in the image for each elementary area of the screen in dependence on the distance of that surface from an image plane and for determining whether any surface is visible at that elementary area, wherein each group of processing elements has an associated cache memory means
30 for storing data defining the various surfaces in the image, and the groups of processing elements are coupled together by a bus means, and wherein data is read from the associated memory means of a first group of processing elements and transmitted via the bus means to the other
35 group of processing elements for determination of depth values until all data from said first associated memory means has been read and reading of data then takes place

- 14 -

from the associated memory means of a second group of processing elements.

Amendments to the claims have been filed as follows

1. A method for shading a three dimensional image for display on a screen comprising the steps of representing each objection in the image as a group of surfaces, for
5 each elementary area of the display projecting a ray into the 3-dimensional image, determining the location of the intersection of the ray with each surface, determining whether any intersected surface is visible at that elementary area, and including the step of designating an
10 object as a light volume, determining the intersection of each ray with the light volume, and shading an elementary area in dependence on the surface visible at that elementary area and its location in relation to the light volume.
- 15 2. A method for shading a three dimensional image for display on a screen comprising the steps of representing each objection in the image as a group of surfaces, determining a bounding volume for each object projected onto the display screen, subdividing the display into a
20 plurality of sub-regions, the size of the sub-regions being dependent on the complexity of the image in each area of the screen, and for each pixel of each sub-region projecting a ray into the image, determining the location of the intersection of the ray with each surface of each
25 object contained with that sub-region, determining whether any intersected surface is visible at that pixel and shading the pixel in dependence on the result of the determination.
- 30 3. A method for shading according to claim 2 including the step of interleaving the determination of ray surface

intersections for rows of pixels in a sub-region with the determination of ray/surface intersections for rows of pixels in other sub-regions.

4. A method for shading a three dimensional image for display on a screen comprising the steps of representing each object in the image as a group of surfaces, subdividing the screen into a plurality of sub-regions, for each pixel of each sub-region projecting a ray into the image, determining the location of the intersection of the ray with each surface of each object contained within that sub-region, determining whether any intersected surface is visible at that pixel, and shading the pixel for display in dependence on the result of the determination, wherein the determination of ray/surface intersections for groups of pixels in one sub-region is interleaved with the determination of ray/surface intersections for groups of pixels in other sub-regions.

5. A method for shading according to claim 4 wherein the groups of pixels comprise lines of pixels in the sub-regions.

6. A method for shading according to claim 4 or 5 wherein groups of pixels from sub-regions in less complex parts of the image are interleaved with groups of pixels from sub-regions in more complex parts of the image.

7. Apparatus for shading a three dimensional image for display on a screen comprising a plurality of groups of processing elements for determining depth values for surfaces in the image for each elementary area of the screen in dependence on the distance of that surface from

an image plane and for determining whether any surface is visible at that elementary area, wherein each group of processing elements has an associated cache memory means for storing data defining the various surfaces in the image, and the groups of processing elements are coupled together by a bus means, and wherein data is read from the associated memory means of a first group of processing elements and transmitted via the bus means to the other group of processing elements for determination of depth values until all data from said first associated memory means has been read and reading of data then takes place from the associated memory means of a second group of processing elements.

Amendments to the claims have been filed as follows

1. A method for shading a three dimensional image for display on a screen comprising the steps of representing each object in the image as a group of surfaces, for each elementary area of the display projecting a ray into the 3-dimensional image, determining the location of the intersection of the ray with each surface, determining whether any intersected surface is visible at that elementary area, and including the step of designating an object as a light volume, determining the intersection of each ray with the light volume, and shading an elementary area in dependence on the surface visible at that elementary area and its location in relation to the light volume.
2. Apparatus for shading a three dimensional image for display on a screen comprising means for supplying data defining a group of infinite surfaces representing each object in the image, means for generating a depth value for each surface for each elementary area of the display in dependence on the distance of that surface from an image plane, means for determining whether any surface is visible at that elementary area, means for designating an object represented by a group of surfaces as a light volume, means for determining a depth value for each surface of the light volume for each elementary area of display in dependence on the distance of that surface from an image plane, and means for shading that elementary area in dependence on the object surface visible at that elementary area and its location in relation to the light volume.
3. A method for shaping a three dimensional image for display on a screen comprising the steps of supplying data defining a group of infinite surfaces representing each object in the image, generating depth values for each surface for each elementary area of the display in dependence on the distance of that surface from an image

- plane, determining whether any surface is visible at that elementary area, designating an object represented by a group of surfaces as a light volume, determining a depth value for each surface of the light volume for each elementary area of the display in dependence on the distance of that surface from an image plane, and shading that elementary area in dependence on the object surface visible at that elementary area and its location in relation to the light volume.
- 5
- 10 4. A method for shading a three dimensional image for display on a screen comprising the steps of representing each object in the image as a group of surfaces, determining a bounding volume for each object projected onto the display screen, subdividing the display into a plurality of sub-regions, the size of the sub-regions being dependent on the complexity of the image in each area of the screen, and for each pixel of each sub-region projecting a ray into the image, determining the location of the intersection of the ray with each surface of each object contained within that sub-region, determining whether any intersected surface is visible at that pixel and shading the pixel in dependence on the result of the determination.
- 15
- 20
- 25 5. A method for shading according to claim 4 including the step of interleaving the determination of ray surface intersections for rows of pixels in a sub-region with the determination of ray/surface intersections for rows of pixels in other sub-regions.
- 30
- 35 6. Apparatus for shading a three dimensional image for display on a screen comprising means for representing each object in the image as a group of surfaces, means for determining a bounding volume for each object projected on to the display screen, means for sub-dividing the display screen into a plurality of sub-regions, the size of the sub-regions being dependent on the complexity of the image in each area of the screen, means for generating a depth value for each surface for each pixel of each sub-region

in dependence of the distance of that surface from an image plane, means for determining whether any surface is visible at that elementary area, and means for shading the elementary area in dependence on the result of
5 determination.

7. Apparatus according to claim 6 including means for interleaving the determination of depth values for surfaces for rows of pixels in one sub-region with the determination of depth values for surfaces for rows of
10 pixels in other sub-regions.

8. A method for shading a three dimensional image for display on a screen comprising the steps of representing each object in the image a group of surfaces, determining a bounding volume for each object to be projected on to
15 the display screen, sub-dividing the display screen into a plurality of sub-regions, the size of the sub-regions being dependent on the complexity of the image, generating a depth-value for each surface for each pixel of each sub-region in dependence on the distance of that surface from
20 an image plane, determining whether any surface is visible at that elementary area, and shading the elementary area in dependence on the result of the determination.

9. A method according to claim 8 including the step of interleaving the determination of depth values for
25 surfaces for rows of pixels in one sub-region with the determination of depth values for surfaces for rows of pixels in other sub-regions.

10. A method for shading a three dimensional image for display on a screen comprising the steps of representing
30 each object in the image as a group of surfaces, subdividing the screen into a plurality of sub-regions, for each elementary area of each sub-region projecting a ray into the image, determining the location of the intersection of the ray with each surface of each object
35 contained within that sub-region, determining whether any intersected surface is visible at that elementary area,

and shading the pixel for display in dependence on the result of the determination, wherein the determination of ray/surface intersections for groups of pixels in one sub-region is interleaved with the determination of
5 ray/surface intersections for groups of pixels in other sub-regions.

11. A method for shading according to claim 10 wherein the groups of pixels comprise lines of pixels in the sub-regions.
- 10 12. A method for shading according to claim 10 or 11 wherein groups of pixels from sub-regions in less complex parts of the image are interleaved with groups of pixels from sub-regions in more complex parts of the image.
- 15 13. Apparatus for shading a three dimensional image for display on a screen comprising means for supplying data defining a group of surfaces representing each object in the image, ~~means for sub-dividing~~ the screen into a plurality of sub-regions, means for each elementary area of each sub-region for generating a depth value for each
20 surface of each object contained within that sub-region in dependence on the distance of that surface from an image plane, means for determining whether any surface is visible at that elementary area, means for shading the pixel for display in dependence on the result of the
25 determination, wherein the determination of depth values for surfaces for groups of pixels in one sub-region is interleaved with the determination of depth values for surfaces for groups of pixels in other sub-regions.
- 30 14. Apparatus for shading according to claim 13 wherein the groups of pixels comprise lines of pixels in the sub-regions.
- 35 15. Apparatus for shading according to claim 13 or 14 wherein groups of pixels in sub-regions in less complex parts of the image are interleaved with groups of pixels in sub-regions in more complex parts of the image.

16. A method for shading a three dimensional image for display on a screen comprising the steps of supplying data defining a group of surfaces representing each object in the image, sub-dividing the screen into a plurality of sub-regions, for each elementary area of each sub-region generating a depth value for each surface of each object contained within that sub-region in dependence on the distance of that surface from an image plane, determining whether any surface is visible at that elementary area, shading the pixel for display in dependence on the result of the determination, and interleaving the determination of depth values for surfaces for groups of pixels in one sub-region with the determination of depth values for surfaces for groups of pixels in other sub-regions.
17. Apparatus for shading a three dimensional image for display on a screen comprising a plurality of groups of processing elements for determining depth values for surfaces in the image for each elementary area of the ~~screen~~ in dependence on the distance of each surface from an image plane and for determining whether any surface is visible at that elementary area, wherein each group of processing elements has an associated cache memory means for storing data defining the various surfaces in the image, and the groups of processing elements are coupled together by a bus means, and wherein data is read from the associated memory means of a first group of processing elements and transmitted via the bus means to other groups of processing elements for determination of depth values until all data from said first associated memory means has been read and wherein reading of data then takes place from the associated memory means of a second group of processing elements.



Application No: GB 9501834.7
Claims searched: 1

Examiner: R F KING
Date of search: 19 July 1995

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK CI (Ed.N): H4T[TBBA,TBBX,TBEA]
Int CI (Ed.6): G06T 15/00, 15/10, 15/50
Other: ONLINE: WPI; H4T Library.

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	EP 0 613 099 A1 (PHILIPS ELECTRONICS) See abstract	1
A	EP 0 610 004 A2 (MATSUSHITA) See abstract	1
A	"Design and Simulation of Opera Lighting and Projection Effects". Dorsey, Sillion, Greenburg. pages 41 - 50, Computer Graphics, Volume 25, Number 4, July 1991	1

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Application No: GB 9501834.7
 Claims searched: 7

Examiner: R. F. King
 Date of search: 28 November 1995

Patents Act 1977
Further Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.N): H4T[TBAD, TBBA, TBEA, TBEC]

Int Cl (Ed.6): G06T 1/20, 1/60, 15/00, 15/10, 15/30, 15/40, 15/50, 15/60

Other: ONLINE: WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
	None	7

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Application No: GB 9501834.7
 Claims searched: 4

Examiner: R. F. King
 Date of search: 27 November 1995

Patents Act 1977
Further Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK CI (Ed.N): H4T[TBEA, TBEC]

Int CI (Ed.6): G06T 15/50, 15/60

Other: ONLINE: WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
Y	GB 2 246 057 A [Rank Cintel] See whole disclosure	4
Y	US 5,305,430 A [Xerox Corp.] See whole disclosure	4

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.



Application No: GB 9501834.7
Claims searched: 2

Examiner: R. F. King
Date of search: 27 November 1995

Patents Act 1977
Further Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.N): H4T[TBEA, TBEC]

Int Cl (Ed.6): G06T 15/50, 15/60

Other: ONLINE: WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
X	GB 2 246 057 A [Rank Cintel] See whole disclosure	2
X	US 5,305,430 A [Xerox Corp.] See whole disclosure	2

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

